



US Patent &amp; Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)Search:  The ACM Digital Library  The Guide

cluster and register and "functional unit" and dispatch and issu



THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used **cluster** and **register** and **functional unit** and **dispatch** and **issue** and **renaming** and **instruction**

Found 13,311 of 142,346

Sort results by

 Save results to a Binder[Try an Advanced Search](#)

Display results

 Search Tips[Try this search in The ACM Guide](#) Open results in a new window

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

**1 Complexity-effective superscalar processors**

Subbarao Palacharla, Norman P. Jouppi, J. E. Smith

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture**, Volume 25 Issue 2Full text available: [pdf\(2.21 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

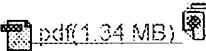
The performance tradeoff between hardware complexity and clock speed is studied. First, a generic superscalar pipeline is defined. Then the specific areas of register renaming, instruction window wakeup and selection logic, and operand bypassing are analyzed. Each is modeled and Spice simulated for feature sizes of 0.8&micro;m, 0.35&micro;m, and 0.18&micro;m. Performance results and trends are expressed in terms of issue width and window size. Our analysis indicates that window wakeu ...

**2 Superscalar architectures: Reducing the complexity of the register file in dynamic superscalar processors**

Rajeev Balasubramonian, Sandhya Dwarkadas, David H. Albonesi

December 2001 **Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture**

Full text available:

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)  
[Publisher Site](#)

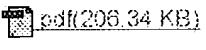
Dynamic superscalar processors execute multiple instructions out-of-order by looking for independent operations within a large window. The number of physical registers within the processor has a direct impact on the size of this window as most in-flight instructions require a new physical register at dispatch. A large multi-ported register file helps improve the instruction-level parallelism (ILP), but may have a detrimental effect on clock speed, especially in future wire-limited technologies. ...

**3 Dynamically managing the communication-parallelism trade-off in future clustered processors**

Rajeev Balasubramonian, Sandhya Dwarkadas, David H. Albonesi

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available:

Additional Information: [full citation](#), [abstract](#), [references](#)

Clustered microarchitectures are an attractive alternative to large monolithic superscalar designs due to their potential for higher clock rates in the face of increasingly wire-delay-constrained process technologies. As increasing transistor counts allow an increase in the number of clusters, thereby allowing more aggressive use of instruction-level parallelism (ILP), the inter-cluster communication increases as data values get spread across a wider

area. As a result of the emergence of this tr ...

#### **4 Reducing wire delay penalty through value prediction**

Joan-Manuel Parcerisa, Antonio González

December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

Full text available:  pdf(148.85 KB)

 ps(401.62 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

 Publisher Site



#### **5 A Complexity-Effective Approach to ALU Bandwidth Enhancement for Instruction-Level Temporal Redundancy**

June 2004 **Proceedings of the 31st annual international symposium on Computer architecture - Volume 00**

Full text available:  pdf(450.01 KB)

 ps

Additional Information: [full citation](#), [abstract](#)

 Publisher Site



Previous proposals for implementing instruction-level temporal redundancy in out-of-order cores have reported a performance degradation of up to 45% in certain applications compared to an execution which does not have any temporal redundancy. An important contributor to this problem is the insufficient number of ALUs for handling the amplified load injected into the core. At the same time, increasing the number of ALUs can increase the complexity of the issue logic, which has been pointed out to be oneo ...

**Keywords:** Complexity-effective design, Instruction Reuse, Temporal Redundancy

#### **6 Energy efficient architectural techniques: Application adaptive energy efficient clustered architectures**

Diana Marculescu

August 2004 **Proceedings of the 2004 international symposium on Low power electronics and design**

Full text available:  pdf(167.43 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As clock frequency and die area increase, achieving energy efficiency, while distributing a low skew, global clock signal becomes increasingly difficult. Challenges imposed by deep-submicron technologies can be alleviated by using a multiple voltage/multiple frequency island design style, or otherwise called, globally asynchronous, locally synchronous (GALS) design paradigm. This paper proposes a clustered architecture that enables application-adaptive energy efficiency through the use of dynami ...

**Keywords:** clustered architectures, dynamic voltage scaling



#### **7 The multicluster architecture: reducing cycle time through partitioning**

Keith I. Farkas, Paul Chow, Norman P. Jouppi, Zvonko Vranesic

December 1997 **Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture**

Full text available:  pdf(1.44 MB)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

 Publisher Site



The multicluster architecture that we introduce offers a decentralized, dynamically-scheduled architecture, in which the register files, dispatch queue, and functional units of the architecture are distributed across multiple clusters, and each cluster is assigned a subset of the architectural registers. The motivation for the multicluster architecture is to reduce the clock cycle time, relative to a single-cluster architecture with the same number of hardware resources, by reducing the size and ...

**Keywords:** decentralized architecture, partitioned architecture, static instruction scheduling, register allocation

**8 A low-complexity issue logic**

Ramon Canal, Antonio González

May 2000 **Proceedings of the 14th international conference on Supercomputing**

Full text available:  pdf(995.88 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

One of the main concerns in today's processor design is the issue logic. Instruction-level parallelism is usually favored by an out-of-order issue mechanism where instructions can issue independently of the program order. The out-of-order scheme yields the best performance but at the same time introduces important hardware costs such as an associative look-up, which might be prohibitive for wide issue processors with large instruction windows. This associative search may slow-down t ...

**Keywords:** in-order issue, instruction issue logic, out-of-order issue, wide-issue superscalar

**9 Overcoming the limitations of conventional vector processors**

Christos Kozyrakis, David Patterson

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available:  pdf(160.23 KB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Despite their superior performance for multimedia applications, vector processors have three limitations that hinder their widespread acceptance. First, the complexity and size of the centralized vector register file limits the number of functional units. Second, precise exceptions for vector instructions are difficult to implement. Third, vector processors require an expensive on-chip memory system that supports high bandwidth at low access latency. This paper introduces CODE, a scalable vector ...

**10 Trace processors**

Eric Rotenberg, Quinn Jacobson, Yiannakis Sazeides, Jim Smith

December 1997 **Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture**

Full text available:  pdf(1.39 MB) 

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Publisher Site

Traces are dynamic instruction sequences constructed and cached by hardware. A microarchitecture organized around traces is presented as a means for efficiently executing many instructions per cycle. Trace processors exploit both control flow and data flow hierarchy to overcome complexity and architectural limitations of conventional superscalar processors by (1) distributing execution resources based on trace boundaries and (2) applying control and data prediction at the trace level rather than ...

**Keywords:** trace processors, multiscalar processors, trace cache, next trace prediction, selective reissuing, context-based value prediction

**11 Microprocessor architecture: A scalable wide-issue clustered VLIW with a**

**reconfigurable interconnect**

Osvaldo Colavini, Davide Rizzo

October 2003 **Proceedings of the international conference on Compilers, architectures and synthesis for embedded systems**

Full text available:  pdf(365.26 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Clustered VLIW architectures have been widely adopted in modern embedded multimedia applications for their ability to exploit high degrees of ILP with reasonable trade-off in complexity and silicon costs. Studies have however shown limited performance scaling for wide-issue machines. In this paper we describe the architecture of a clustered VLIW with a runtime reconfigurable inter-cluster bus suitable to address such scalability problem. The architecture is aimed at kernel loops acceleration thr ...

**Keywords:** IDCT, clustered VLIW, modulo scheduling, reconfigurable co-processor (RCP)

**12 Reducing the complexity of the issue logic**

Ramon Canal, Antonio González

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Full text available:  pdf(140.22 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

*The issue logic of dynamically scheduled superscalar processors is one of their most complex and power-consuming parts. In this paper we present alternative issue-logic designs that are much simpler than the traditional scheme while they retain most of its ability to exploit ILP. These alternative schemes are based on the observation that most values produced by a program are used by very few instructions, and the latencies of most operation are deterministic.*

**Keywords:** complexity-effective design, instruction issue logic, out-of-order issue, wide-issue superscalar

**13 Optimization of high-performance superscalar architectures for energy efficiency**

V. Zyuban, P. Kogge

August 2000 **Proceedings of the 2000 international symposium on Low power electronics and design**

Full text available:  pdf(196.55 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years reducing power has become a critical design goal for high-performance microprocessors. This work attempts to bring the power issue to the earliest phase of high-performance microprocessor development. We propose a methodology for power-optimization at the micro-architectural level. First, major targets for power reduction are identified within superscalar microarchitecture, then an optimization of a superscalar micro-architecture is performed that generates a set of ...

**14 Low-power: Low-complexity reorder buffer architecture**

Gurhan Kucuk, Dmitry Ponomarev, Kanad Ghose

June 2002 **Proceedings of the 16th international conference on Supercomputing**

Full text available:  pdf(120.97 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In some of today's superscalar processors (e.g. the Pentium III), the result repositories are implemented as the Reorder Buffer (ROB) slots. In such designs, the ROB is a complex multi-ported structure that occupies a significant portion of the die area and dissipates a non-trivial fraction of the total chip power, as much as 27% according to some estimates. In addition, an access to such ROB typically takes more than one cycle, impacting the IPC adversely. We propose a low-complexity and low-power ...

**Keywords:** low-complexity datapath, low-power design, reorder buffer

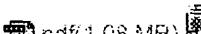
**15**

An instruction set and microarchitecture for instruction level distributed processing

Ho-Seop Kim, James E. Smith

May 2002 **ACM SIGARCH Computer Architecture News**, Volume 30 Issue 2

Full text available:



[pdf\(1.08 MB\)](#)



Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

[Publisher Site](#)

An instruction set architecture (ISA) suitable for future microprocessor design constraints is proposed. The ISA has hierarchical register files with a small number of accumulators at the top. The instruction stream is divided into chains of dependent instructions (strands) where intra-strand dependences are passed through the accumulator. The general-purpose register file is used for communication between strands and for holding global values that have many consumers. A microarchitecture to supp ...

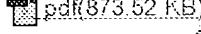
- 16 [Instruction distribution heuristics for quad-cluster, dynamically-scheduled, superscalar processors](#)



Amirali Baniasadi, Andreas Moshovos

December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

Full text available:



[pdf\(873.52 KB\)](#)



Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

[Publisher Site](#)

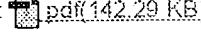
- 17 [Banked multiported register files for high-frequency superscalar microprocessors](#)



Jessica H. Tseng, Krste Asanović

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available:



[pdf\(142.29 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Multiported register files are a critical component of high-performance superscalar microprocessors. Conventional multiported structures can consume significant power and die area. We examine the designs of banked multiported register files that employ multiple interleaved banks of fewer ported register cells to reduce power and area. Banked register files designs have been shown to provide sufficient bandwidth for a superscalar machine, but previous designs had complex control structures that w ...

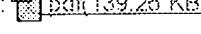
- 18 [Improving dynamic cluster assignment for clustered trace cache processors](#)



Ravi Bhargava, Lizy K. John

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available:



[pdf\(139.26 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

This work examines dynamic cluster assignment for a clustered trace cache processor (CTCP). Previously proposed cluster assignment techniques run into unique problems as issue width and cluster count increase. Realistic design conditions, such as variable data forwarding latencies between clusters and a heavily partitioned instruction window, increase the degree of difficulty for effective cluster assignment. In this work, the trace cache and fill unit are used to perform dynamic cluster assignme ...

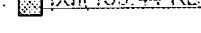
- 19 [Clustered microarchitectures: Cluster prefetch: tolerating on-chip wire delays in clustered microarchitectures](#)



Rajeev Balasubramonian

June 2004 **Proceedings of the 18th annual international conference on Supercomputing**

Full text available:



[pdf\(153.44 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The growing dominance of wire delays at future technology points renders a microprocessor communication-bound. Clustered microarchitectures allow most dependence chains to execute without being affected by long on-chip wire latencies. They also allow faster clock speeds and reduce design complexity, thereby emerging as a popular design choice for

future microprocessors. However, a centralized data cache threatens to be the primary bottle-neck in highly clustered systems. The paper attempts to id ...

**Keywords:** clustered microarchitectures, communication-bound processors, data prefetch, distributed caches, effective address and memory dependence prediction, processor

**20 Putting the fill unit to work: dynamic optimizations for trace cache microprocessors**

Daniel Holmes Friendly, Sanjay Jeram Patel, Yale N. Patt

November 1998 **Proceedings of the 31st annual ACM/IEEE international symposium on Microarchitecture**

Full text available:  [pdf \(1.90 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



Results 1 - 20 of 200

Result page: **1** 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)